Oral History of Eli Harari

Interviewed by: George Alexy and Jim Porter

Recorded: June 15, 2011
Mountain View, California

CHM Reference number: X6066.2011

Porter: Well, we're going to start by having this very interesting interview with Eli Harari. We have two interviewers. I'm Jim Porter, I have a background in the disk area for 23 years, and I publish market studies on the disk drive industry, and I can say the first disk drive I ever saw was the first disk drive, when I went to work for the company that got that one back in the '50s. But let me introduce George Alexy, who has a background which is different from mine, and who will start this interview by going into the background with Eli.

Alexy: And I'm George Alexy, member of the Computer History Museum's Special Interest Group in Semiconductor Memories, so I'm going to be coming at the interview from a semiconductor perspective. We're very pleased to have Eli Harari here with us today. Now retired, CEO, Founder and President of a large company that is well known in the valley and around the world, SanDisk, also an entrepreneur and technologist, so welcome, Eli, we're very pleased to have you here for this interview.

Harari: Thank you, it's a pleasure, pleasure to be here.

Alexy: I thought we might start with some background from you, such as when you were born, where you grew up; the early days.

Harari: Yes, so my parents came from Poland-- I was born in Israel, at that time, 1945, June, 1945, it was still Palestine, controlled by the British. I'm told that when I was eight days old, my parents took me back home from the hospital in a pram and the British-- there was a curfew, and the British soldiers stopped my parents with the pram and searched me in case I had something, and this was eight days old. Anyway, I grew up in Israel until age 13, and then went to England to boarding school. My parents wanted me to become an English gentleman, and at 18, I entered the Israeli military, compulsory service, slightly over two years in the air force. I was a technical clerk, and then started my education, first physics at Manchester University, England.

Alexy: So was this, the air force, your first introduction to technology and electronics?

Harari:You could say that, yes.

Alexy: So, as you went to university, what universities did you attend, and what was it that got you interested in the course of study that you pursued. What was the driving factor that captured your imagination and lit that spark?

Harari: Well, without any doubt, it was Kennedy's-- Jack Kennedy's speech about 1961, you know, landing a man on the moon, and returning him safely, before the end of the decade. So I was 16 at that

time. In '57, when I was 12, was the Sputnik, of course, so we grew up in an age where, you know, there was no limit, I mean, the sky's the limit, space, and I was really inspired by Kennedy's vision, and in fact, arrived in the United States to do my PhD work here, just a month after Neil Armstrong actually stepped on the moon.

Alexy: Fascinating.

Harari: Yes.

Alexy: So where was your first university, you said, in England, was it?

Harari: Yes, I attended one year at the Hebrew University in Jerusalem, studying math and physics, then transferred to Manchester for physics, graduated in 1969, with a physics degree, and then came to the United States, Princeton University, Princeton, New Jersey. Princeton, I did my PhD work in what is called solid state sciences, which is really a branch of physics, applied physics, and graduated in '73.

Alexy: So is this what got you involved in semiconductors, the applied physics and solid state physics?

Harari: Yes, yes.

Alexy: So what types of things were you working on at the university?

Harari: The university, this was the height of the cold war, and you know, early satellites, communication satellites, spy satellites, were failing in space, so I was funded, my PhD work was funded by the Office of Naval Research, and the navy wanted to understand why satellites were failing after a very short time in space, earth orbits. After three months, they would fail, the electronics would stop working. So my thesis work was really-- my research work was to try to understand the failure mechanism in radiation hardened devices, and it turned out, to my great luck, that I was working on the right materials, which was SiO2, silicon dioxide, films of SiO2, and aluminum oxide, thin films of that. And the physics that I developed, as an experimental physicist, turned out to be the physics that influenced very much, later on, my career in non-volatile memories, flash memory.

Alexy: So as you finished at Princeton, what was your first job out of university?

Harari: I came to the west coast, I had a job offer from Hughes Aircraft and Rockwell, both to work on space radiation effects, basically semiconductors in space. I worked at Hughes. Hughes was a private non-profit company, it was owned by Howard Hughes, it had phenomenal labs. You could do just about

anything you wanted, you know, your direcTV is a consequence of the satellites that were launched in those days. And at Hughes, I was given a free hand. I could do, basically, whatever I wanted, so long as part of it involved space radiation. And the challenge I set for myself was really to try to develop an erasable EPROM. So around 1970, Dov Frohman, at Intel, invented the floating-gate EPROM device. The floating-gate EPROM device was a semiconductor memory that was encapsulated in a package that had a quartz window. It could be programmed with a program that the microprocessor, basically dictated what the microprocessor would do, and it was erasable by ultraviolet light through that window, which took about 15 minutes, so you had to take this part out of the board, put it under a lamp, and erase it. EPROM, Frohman's EPROM were a very pivotal component in the microprocessor revolution, basically it enabled, you know, a general purpose computing machine to have different instructions, through the codes on that EPROM. What I was trying to do was to develop an EEPROM, and electrically erasable EPROM, as opposed to UV erasable, and that I did at Hughes, and in fact, came up with the idea that if you thin down the gate oxide of that floating gate, the Frohman-Bentchkowsky device from 1,000 angstroms, to 100 angstroms, then that 100 angstroms would, in fact, allow you to program and erase electrically. This is called, Fowler Nordheim Tunneling, after two gentlemen, British gentlemen, that discovered that phenomenon.

Alexy: And you applied it to this particular problem.

Harari: Yes, right. So the problem, at that time, and this actually has tremendous relevance to the flash memory industry, tremendous impact. The problem was that there was very little known, in those days, this was 1976, '77, you know, 35 years ago, very little known about thin films of SiO2. SiO2 was used in MOS transistors, but, as I said before, it was about 1,000 angstroms thick. The thinking was that if you thinned it too much, let's say, to 100 angstroms, it would, basically, be very leaky, pin holes and breakdowns and so on. In order to make an EEPROM device, that used that very thin oxide, I had to basically develop the whole technology of thin, very thin SiO2.

Alexy: Is it a manufacturing technology taken to production?

Harari: Yes, so at Hughes, again, I had this tremendous freedom, I could do any-- we did tens of thousands of samples over a two year period, and developed a very, very sophisticated technology in metrics, to characterize these very thin films. And what I found out was that these films actually were very reproducible, that you could make them with very high quality, very high integrity, that they would conduct very-- conduct electrons both for programming and erase, between the substrate and the floating gate, but that they eventually broke down because what would happen under very high electric fields, that were required to program and erase, the oxide would, over time, degrade, it could only pass a certain amount of charge, electrons, back and forth, and in the process, you were creating traps in the oxide in the thin oxide, electrons started getting trapped in that oxide, and it would, basically, eventually fail. But before it failed, my calculations showed it could achieve write-erase to over a million times, which was more than adequate.

Alexy: That's what I remember, because one of the key marketing elements of these devices was the number of read-write-erase cycles that you could perform on them, while-- and still maintain the integrity of the device, which, in most applications, was not that frequent at all.

Harari: Yes, right.

Alexy: So Hughes was really a very good place to explore and develop. Were there any good stories of people you worked with and things that went really well, and things that didn't go so well, and--

Harari: Yes, I have some very good-- my boss, Tom Toombs, who was the manager of the research center, also a Princeton PhD, gave me a free hand and all the support and I was able to do a phenomenal amount of data gathering and experimentation that gave us a very, very clear understanding that this actually was a very viable-- these thin films of silicon dioxide were very viable, and we very quickly went into reduction to practice. We-- I don't mean just by filing patents, but more importantly, developing the industry's first floating gate EEPROM chip, it was an eight kilobit CMOS EEPROM.

Alexy: And did Hughes put that device go into production?

Harari: Yes, it went into production around 1980/81, and in addition to that, we developed the first NOVRAM, which was a static RAM chip that was backed by a nonvolatile cell, so that when you-- when the power was turned off, whatever state the static RAM six transistor cell was in, would be programmed into the flash-- it was not flash, into the EEPROM. So these were the, kind of precursors for the industry, adopting EEPROM. Intel came out with their own EEPROM, Seeq was another company that developed that, and basically the EEPROM industry was developed. Dov Frohman, who was a key member of Intel's team, saw, you know, saw my publication. I published very extensively, the work that I did at Hughes, on these thin oxides, in the Journal of Applied Physics, in 1978. He saw that, he jumped on the next plane, came to Hughes and recruited me to come to Intel. And that's where I-- that's how I came to the valley in 1979.

Alexy: Regarding your experience at Intel, you were there for how many years?

Harari: I was there two years.

Alexy: And you worked with Dov Frohman at the time?

Harari: No, Dov had gone to Israel, to start Intel Israel. I was in the-- I was manager of the Santa Clara Technology Development Group. Intel, at that time, had three technology development groups, one was

Santa Clara, one was Portland, Oregon, and one was-- what do you call it, Livermore-- not Livermore, but-- well it's Livermore.

Harari: Yes. Livermore was responsible for HMOS, static RAM, Portland was responsible for DRAM, I was responsible for nonvolatile memories, and basic technology. Basic technology was basically developing the first steppers, the first dry etchers, the first deposited oxides for the next generation of process technology.

Alexy: Who were some of the people you worked with in the Intel days?

Harari: So my boss was Gerry Parker, he was the head of all technology at Intel. And you know, there were the major people that helped develop the technology, Ted Hoff, was-- used our pilot line in Santa Clara. Ted Hoff once told me, this was 1980, and this is, again, relevant to my career, he once told me, "If you can develop a memory technology that has 10x lower cost than the existing technology, the entire hierarchy of memories will move away, basically will part way to allow your technology to come into that hierarchy," and that stuck with me really, until today, he was right on the money. So the effect of 10x, in cost reduction is critical, and that's the-- an early lesson from Ted Hoff.

Alexy: While you were at Intel, what products were developed at that stage? How had this technology evolved from your early Hughes days, up through the state of the art at Intel?

Harari: So Intel was focused very much on EPROM, that was still the big money maker. My team developed the first stepper based dry etch at 64 kilobit EPROM at Intel, which, for three years, carried Intel, it was the most profitable product at Intel. And we were working on relatively low density EEPROM, 16 kilobit or so, so that was EEPROM really never made it, as a major mainstream technology. It was a two transistor cell, and not particularly low cost.

Alexy: So the belief factor wasn't there yet, within Intel for the technology.

Harari: No, I think the breakthrough for EEPROM came, really in 1984, with Masuoka at Toshiba, and at Exel, I think, Mukherjee and Chang, at Exel, which was a small startup, that said, look, we can simplify this EEPROM, take out the select transistor. All we need is to erase the whole chip electrically instead of with a UV lamp, that's already a huge step forward. Let's simplify it and that would make it cheaper. That was the invention of flash memory. Flash, basically, it's really an EEPROM-- it was actually called, flash EEPROM. It was still an EEPROM, but you eliminated the select transistor and you bought yourself a much simpler structure, basically one transistor that had thin oxide, in the case of the first Masuoka flash device, actually it never worked, it was a triple poly structure. The Exel structure actually was the foundation of what Intel later called, ETOX flash, NOR flash.

Alexy: So up until that time, the EE squared was erasable and rewritable and they moved from the bit level to a block level, or a full chip level?

Harari: The first flash was a full chip level. Seeq then introduced sector-- basically a page level, you could erase either the whole chip, or one sector.

Alexy: You were at Intel for, two years, was it?

Harari: Yes.

Alexy: Where did you go after that, and what was it that got you excited about going to where it was you went?

Harari: And so after Intel, I went to Synertek, it was a Honeywell company another startup company whose major customer was Apple, the first Macintosh, the Apple II, used a 6509, I think, 8 bit processor from Synertek. And Atari, Atari was the ROM cartridges. Then Atari was acquired by Warner Brothers, and eventually the company went out of business, Nintendo drove them out of-- basically. So at Synertek, I was VP of Technology, and then VP of Operations and Technology. I was there for about two years.

Alexy: Was this your first startup experience?

Harari: Well, Synertek, I was not-- no, Robert Schreiner was the founder, and he had left, and I left when he left. No, my first startup was after that in 1983, I started a company called Wafer Scale Integration, WSI, and the original business plan, original business plan was to develop basically mass storage systems using EEPROM. This was '83, so at that time there was no flash yet. And I figured I would find a way to reduce-- lower the cost of EEPROM, but at Wafer Scale, we actually ended up doing a completely different business plan. We never did what the original business plan--

Alexy: Why was that?

Harari: Because it was too ambitious, the plan to develop mass storage with EEPROM was too far ahead of its time, if you will.

Alexy: So the technology was not quite ready for that level of product?

Harari: Not quite ready, right, and we couldn't find anybody to fund that development so we went instead and developed CMOS EPROM and programmable logic devices. I was at Wafer Scale, I was for the first two years, I was CEO, then for the next two years, I became Chairman and CTO. I brought in a CEO, and after four years at WSI, I left the company and started SanDisk.

Alexy: How was it at Wafer Scale, in terms of your first experience as the founder, in building a company? What takeaways did you take from that, as you moved forward in your career?

Harari: I took away many, many lessons, unfortunately very painful lessons. I have lots of scars until today, but they were-- I was very inexperienced, I made a very large number of mistakes, and we would really have to be here for the next several hours, to go over all the mistakes I made, but every one of those mistakes turned out to be very valuable later on, because I usually don't repeat the same mistakes.

Alexy: One of the interesting aspects of these interviews is to really gain the wisdom from those that have walked the path and gone through these experiences, on to greatness and building great companies. Are there any of those thoughts you can share with us? They are always valuable to the rest of us in the industry.

Harari: Sure. Well, certainly I consider myself a second chance guy, because my first chance at a startup, I consider a failure, and probably most investors would consider a failure, even though, eventually, the company was sold to STmicro. You know, the first mistake is, you think you can do anything. You can do, you know, you can do memories and programmable logic, and you know, basically, and of course you can't. In a startup, if you want to succeed, you need to focus, you really need to focus. So the key, you need to have a laser sharp focus and the discipline to not stray left and right. You need to hire really a very, very talented, dedicated, passionate and complimentary team. Complimentary, that means, their skills are different from yours, together you make the whole, but that share your vision. The vision must be very simple, very clear, that 30 second elevator speech needs to be really believable and achievable. You need to have a very clear strategy, how to get there, you need to have enough cash in the bank, because it always takes much, much longer than you think to get there. So if you have the right strategy, the right focus, the right people, and the right cash to get there, then the rest is technology, really, what is your unfair competitive advantage, because you know, as soon as you become anything in the market, your competitors will do their best to kill you. There's no-- this is not a-- this is a pretty hostile environment for startups, as we found out. So we didn't have-- under my CEO leadership, I was a technologist, I was not a CEO, I was a CEO that was not really a good CEO, and was not focused. We were doing too many things, we had a board of directors that was really two separate boards. One thought they were investing in a memory technology company, one in a programmable logic company. We spent too much money, you could say we never had enough money, but we didn't give it the respect that you need. So at the end of the day, you have to take risks, a startup company is all about risk taking. There is no such thing as not taking risk. If you don't-- if you don't take risk, you know, the way to fail in a startup is to take too much risk, or stupid risk, bet your company kind of risks unnecessarily so, or to not take enough risk. The key is,

what's the right kind of risk, and where do you put your bets, and where do you bet your company, if you do, and do it as little as possible, and try not to be betting your company in situations where you have no control of the forces that will determine the outcome. So--

Alexy: So at Wafer Scale, what were some of the forces that you realized you didn't have control of, that ultimately were critical?

Harari: Well, we had a divided company that, you know, not only was the board divided, we had a small company with 100 and some employees, half of whom were doing memory and half of whom were doing logic, and the rationale was, this is programmable logic, you need to have the memory technology just pay the bills, but to be good in the memory technology, you have to be darn good, and to be good in the programmable logic, you need to have fantastic software, those were two different--

Alexy: Were you a fabless company at that time?

Harari: Yes, we were a fabless--

Alexy: So how did that work? Did you have to do technology development, or were you working through the technologies available, because the foundry industry was just getting started at that point, wasn't it?

Harari: Yes, 1983 was the good times for LSI Logic, for VTI, VLSI Technology, they were riding high. Yes, and we thought, hey, we have the programmable element that we knew how to make EPROM, E-square PROM, we were going to program it instead of through a mask, a dedicated mask. So the idea was good, but it was, again, overly ambitious and we didn't have enough, even though we raised quite a lot of money, I think, 40, 50 million dollars, over four years, that was not enough.

Alexy: Were you developing your own fab, or were you working with foundry partners to implement the technologies you needed?

Harari: We worked with foundry partners, mostly Sharp in Japan. Sharp was a very good partner.

Alexy: Eli, thank you. Interesting strategy. At this point I want to start to bring Jim into the discussion, because I think we're moving from the semiconductor phase, into the founding of SunDisk, which ultimately became SanDisk. What was the story behind that change?

Harari: So in '88, when I started SunDisk, I was looking for a name. And, by the way, just to close on WSI, I left WSI really at the kind of gentle push out, if you will, of the board, because the CEO that I

brought in, and I, were not getting along well together, and that's another lesson for entrepreneurs. That switch, entrepreneur stepping aside to bring a professional manager, is deadly dangerous. It can kill the spirit of a company, and kill the, you know--

Alexy: The heart--

Harari: The heart, the soul of the machine, if you will, and that's basically what happened at WSI. So when I started SunDisk, which was, I was 43 years old, you could say a failure, and you know, 43 is not a young age to start a company, but I left on February 28th, 1988, and started my new company March 1st, the next day, 1988, and I was optimistic as hell, I was just raring to go, and I had-- before leaving the company, I disclosed to the CEO and to the board that I felt that flash was coming around the horizon and that it was beginning to look like the right time to develop a mass storage based on flash memory. The company and the board said, be our guest, you know, we-- there's no way we can do that as well as doing programmable logic and EPROM, so it was an amicable separation, but parting.— So to get back to SunDisk, my daughter, who was 15 at that time, we were looking for names, and she came up with the name, SunDisk, S-U-N-D-I-S-K. She said, it's just a sunny, uppity name. I remember, we were driving around Fremont and she was in the car, and we said, I said immediately, this sounds a really great name. And that was the name that we had until 1995, for seven years. Before we went public, Sun Microsystems had been on our case. Once we became relatively well known, they said, look, we may have a disk drive division, Sun Disk, or Sun Soft. They actually had a Sun Soft division, so you have to change your name. And they started suing us in different countries for trademark. So I said, look, we've got some real wars out there, and this is one that we don't need to have. I went to Sun, met their CFO, I forgot his name, but a very decent guy, and said, look, we are prepared to change our name to S-A-N, SanDisk. SanDisk actually is sand. Sand, silicon disk, or SansDisk, no disk, so it was actually good, it sounded the same, and in Japan, you didn't even have to change the letters. In Japanese, SunDisk and SanDisk are the same. But I said, but it's going to cost us. You need to help us. So they agreed to pay us, basically the full cost. They gave us five years to phase out in our products, and basically paid the bill for all the copyrights and all the literature, everything that cost us- very, very classy. It would have cost them, and us, a hell of a lot more to fight it in court, of course, so this was a very, very good example of common sense prevailing.

Alexy: Picking your battles carefully.

Harari: Picking your battles carefully, thank you.

Alexy: So, Jim, you had some questions particularly around this.

Porter: In '88, you were cofounder of SanDisk, and became president and CEO. My understanding was, shipments didn't really start until '93. Is that correct?

Harari: We actually shipped-- first, let me say that my cofounders, who were with me, I mean, all the way, my second founder was Jack Yuan, who was our process engineer, great process engineer, just recently retired, Sanjay Mehrotra, was our third founder, he was our memory design/product engineering guy, and he is today the CEO, so when I stepped down, he took over, 23 years later, and that's really very, very rare for the three founders to be around from the early days, through all the growth phases of the company. And the fourth member is the one that I think you know, Bob Norman, who was our system architect and really a great guy, really a great guy. So we started shipping our first product in 1991, the customer was IBM, Boca Raton. IBM at Boca Raton, at that time, was king of the PC world, Entry Systems Division. They were developing the industry's first pen computer. They called it the ThinkPad, because it was a pad, there was no keyboard. It was an eight pound machine, it was 8 by 11, and they needed-- they were using a two and a half inch 20 megabyte drive from Connor Peripherals, and it was failing because you would go for inventory in the factory, you would drop it, the disk drive would fail. So IBM sent out a request for proposals from the industry, to develop a 20 megabyte Flash, two and a half inch, plug and play compatible ATA drive that would replace the 20 megabyte Connor two and a half inch hard disk drive. And they said, whoever does it, the RFP was for 10,000 units, 10,000 20 megabyte Flash drives. For us this was unbelievable, big. So we bid on it, so did Western Digital, who was also an investor in SanDisk, but we competed for that bid, and I believe TI, also bid on that thing. Cut a long story short, we won that bid, which was amazing, because we were, like, 40 people in the company, at the time that we won it, and at a certain point in time, I know for a fact, that IBM had more people working at Boca Raton on helping us than we had at SanDisk. But to cut a long story short, when we were ready, and we shipped, and product was-- this was basically the product, it had our own controller with off the shelf components, and it had cards, memory cards that used our own flash chip, so each card was 10 megabyte, we had two cards, 20 megabyte, and you could-- and this was a standard Connor interface, and this was a Cirrus Logic interface chip, so you could just-- and firmware, of course, you unplugged the Connor, you plugged this, and it was solid state. To the computer, it emulated a hard disk drive instruction set, everything of the ATA instruction set. The computer thought it was talking to a disk drive. Of course, there was no rotating anything, and it was really a major, major breakthrough, because it took flash that was very unreliable, and made it totally-- its weaknesses totally transparent to IBM. However, the story of this, is that when we shipped the first units to IBM, engineering samples, they would fail after about one to two seconds of their diagnostic software. They would run it. So I talked to the key IBM reliability engineer, and I said, "Well what gives?" He said, "Well, look, don't worry, every time we get a new disk drive, and we throw this very, very-- we throw the book at the drive, it usually fails quite quickly." I said, "Well, when would you consider this to be a reliable device?" He said, "Well, if I can run it over the weekend-- if I can run several units over the weekend, with no failure, with our diagnostic software, then you have a good product." So I said, "Well, give me a call when we reach that point," and it took six months before I got that call.

Porter: Well, back in the early days of SanDisk, which we've been discussing, what was your initial evaluation of the opportunity in the storage markets, and what were all those kinds of markets for storage devices such as you were putting together?

Harari: The original business plan, which I wrote in 1988, second half of 1988, saw just about every major opportunity that we eventually ended up delivering on. In fact, in January, 1990, just, really 15 months after the company was started, I gave an invited talk at the IEEE, Santa Clara Chapter, in Santa Clara University, and I still have that abstract from that meeting. And in that, I talk about use of flash memory in cell phones, in digital cameras, in portable fax machines, which was a phone, before there was-- there was no e-mail, of course, in those days, so we basically saw the applications, of course the industrial applications, the military applications, so-- but the issues were, we'll get back into that, but the issue was that we recognized, because of my background as a device physicist, I recognized that flash was unreliable. Flash could not-- it could achieve a million cycles, but quite a few bits would fail before the distribution, the vast distribution would reach a million cycles. So how do you guarantee-- in fact, the spec for E-squared PROM was typically 10,000 cycles, and when I went and talked to customers, potential customers in the early days of SanDisk, they would say, 10,000 cycles is a nonstarter. If you don't have a million write-erase cycles, don't bother. So how to get from 10,000, which by itself, was very difficult to do, to a million, with zero failures, was a challenge, number one. The second challenge was, how do you make it low cost, because we were, you know, our first product was 50 dollars per megabyte, that's 1,000 dollars for that 20 megabyte. Not too many people would afford that kind of-- so we had to make-- to take a technology that had tremendous potential, but was like a frog that needed to become a prince. How do you take this technology that has the potential, statistically, to reach there, but you know-- I can't tell you, the next time I write into that memory, if any bit will fail, and which one will. And the only solution was what we call “system-Flash”. Basically it said, the physics is the physics. The physics of failure, we understood. We knew it, and statistically, you could predict very easily how many bits would fail. But you know, standard techniques in those days were hardware redundancy, you had memory arrays with extra rows and extra columns, if a bit failed, you substituted it with an extra row or an extra column. You tested it, it worked, you shipped the product. This was true for an EPROM, or even for an EEPROM. It was true for DRAM, because DRAM didn't have a wear out mechanism, static RAM didn't have a wear out mechanism. Flash had wear out, and how do you overcome that wear out? The only solution we concluded, was a system solution, a holistic solution that said, you need a controller that works closed loop with that flash memory. The flash memory itself needs to be architected to emulate a disk drive with a sector, with a header, a header accessible by the controller, and data in the sector itself, accessible by the user. The controller had to work closed-loop continuously with the flash memory, detecting bits before they were failing, before they were about to fail, or, if they failed, replace them, map them out. If too many bits fail, map out an entire sector, map out entire region of the memory. Write data across the board, level out the wear out. Don't physically write always to the same physical cells, drive them, you know-- so we had things like, developed hot count, which was really every sector, every time we erased it, we incremented by one, the number of times it went through a write-erase. Every sector had different voltage applied to it during write and erase, to minimize the stress on that oxide. If you look at an Intel chip in those days, Intel would apply the maximum end of life voltage. It, for the first pulse. We said, no, you know, in our case, the end of life voltage, to get to a million cycles, we need to internally generate, maybe 24 volts. But the first cycles, we need only apply 12, 13 volts. So the wisdom in those-- the standard wisdom in those days is, you know, guarantee 24 volts and program every time at 24 volts. We said, well if you want to get to a million cycles, program at the minimum voltage that you need to, and let's say 13 volts, well that voltage was stored in the sector, in the header, and next time we came to that sector, we applied 13 volts, not 24. And that gave us the million cycles. The hot count, the wear out leveling, the

programmable voltages, all of those things that were completely revolutionary could not be done at a chip level, but this is what the controller could do. And that was how we took flash from really what we call a code store, storing code, which is rarely ever changed, maybe once or twice, to storing data and content and software and images and music and so on.

Porter: Well, your company had a rather interesting interrelationship with some disk drive manufacturers. You had investments from companies like Seagate, and Western Digital, for example. What were those relationships like, and why were those companies so interested in your company?

Harari: 1988, when we started the company, my Chairman was Irwin Federman, who was also on the board of Western Digital. So he introduced me to Western Digital at the highest levels, including the technology team, but also the business side. And they were sold, in 1988, on solid state drive, on flash being the future technology. They were very, very early. Kathy Braun, I'm sure you know-- you remember her, and Carl Lofgren, very, very good guys, and they invested two million dollars in the company. Unfortunately for them-- and we worked very closely together, and they brought some of the controller concepts, the ATA, they were pioneers in ATA. The problem for them was that their disk drive business was-- they went into a trough, and they just could not continue to invest in our company, and they basically pulled out, sold their stock. In 1992, when we were four years old, Seagate came to us, and they actually-- Garrett Garretson, you probably remember that name from the past, was very intrigued by solid state disks, and came to us and spent some time with me, and then Al Shugart, Seagate's founder and CEO at the time, came and said, look, you know, this looks like very interesting technology, but we don't know where it's going to go, maybe it's too early. I'd like to make a personal investment in your company. You know, just as an investor. I said, "Well, let me talk to my board," and I came back to him, I said, "Look, Al, we'd be really greatly honored to have your investment, but what we really want is for Seagate to have a strategic partnership with us. So if you can get that, then it will be, you know, that's really what we'd like to do." So they took several additional months, and eventually, they concluded that they wanted to acquire the company. This was 1992. Well I wasn't interested in selling the company, so we started negotiating, Al and I, and we were-- I was not prepared to sell more than 25 percent of the company and he wanted 35 percent, which I felt was as much-- as good as selling the company. Cut a long story short, he accepted 25 percent, they put 30 million dollars into the company, and we had a very, very good relationship for quite a number of years. They were clearly ahead of the time, as far as solid state disks, really just now coming into being, because the cost wasn't right, it was way too expensive. And even now it's just kind of borderline crossing into affordable. So Al joined our board, stayed with our board until very close to his passing away, was a great, great board member and Seagate was a very good partner, very good partner. The original plan with Seagate was that they would be our sales channel. We would manufacture-- do the development, do the manufacturing of these 20 megabyte drives, and they would sell it to their customers. That turned out to be-- that almost killed us. Even though it was with the best intentions, it really was with the best intentions. They were not trying to kill us, but they couldn't get their sales guys-- could not get their sales guys to sell the product, because it required a very, very long design-in cycle, it was a brand new technology, the sales force didn't know anything about flash memory. So a year into this thing, we were selling nothing, you know, diddly, so I went to Al, and his CFO, and VP of Worldwide Marketing and Sales, and I said, "Look, guys, you're killing the company." We're spiffing,

you know, special promotion incentives to our sales force, every design win is 20,000 dollars, sounds great, but if they don't meet the quota of 40 million dollars a quarter of disk drive sales, they're out. So they go for the quota. So this is not working. So Al says, "Okay, what do you want us to do?" I said, "Well, give us these specific seven guys, your sales guys, because we know now, after a year, that these guys are the fanatic guys that really love flash memory. And get out of our way. Let us handle the sales, just forget about what we have in the contract, we'll take over. Make sure that Hamilton, Avnet, Arrow, all the industrial distributors are still with us, because we were tiny, and help us at COMDEX and so on." They did all of that, all of it. I mean, it was just wonderful the way they were very mature about it, they accepted it, and that really helped us. That's when we really started selling. So 1993 was about the time.

Porter: Okay, so going back to the markets how has the original concept for SanDisk products and markets evolved with the evolution of the market for data storage products, as we have all kinds of applications now that are being used with flash memory, how has your concept of the market opportunity for the company evolved to take advantage of all that?

Harari: So very early on, we-- as I said, from day one, we knew that the cost had to come down dramatically, that there was no way we were going to build a mass market with the cost of, you know, 50 dollars per megabyte or even one dollar per megabyte.

Porter: It's been suggested that price is the most important technology product.

Harari: Yes, and Moore's law came in very, very handy in that regard, so we definitely practiced Moore's law to the hilt, starting in 1991, with our first product, used our own four megabit flash chip, it was the only four megabit Flash chip in the market. And we went from four megabit in 1991, to 64 gigabit chip today, 64 gigabit, so we went through 14 generations, 14 doublings, 2 to the 14, and that allowed us to bring the cost of flash down, over that period, by about 30,000 times, 30,000 times cumulative cost reduction, which means that if you were to buy a 3,000 dollar PC in 1991, 3,000 dollars, you would pay for it today, ten cents, that's a 30,000X factor. I remember Ted Hoff saying, "If you reduce the price by a factor of ten," so here you have ten times ten, times ten, times ten, times three, and we'll get the next ten. So every time, we were able to bring the cost down by a factor of ten, we would invade a new market, and we would disrupt an existing market. The first one, of course, was the film, and digital film replacing silver halide. Silver halide, think about it, was, you know, $3.99 a roll of film, and phenomenal quality, and here we are, out at our first card, which we made for Kodak, this was a two megabyte card that we sold Kodak for 100 bucks, 50 dollars per megabyte. So we knew it has to come down. First of all, it had to come down in size. This was way too big for anything portable, secondly, we said, this is a very low power technology, it is nonvolatile, you don't need battery. You can store information indefinitely, so it's ideal for storage of information, content, mass storage in anything portable. Anything portable needs to be very small. So our first product we condensed all of this, into this, which is the PCMCIA card, which we were founding members. This had, already all the controller inside, as well as the flash memory, and it had a standard interface. We even went to the disk drive guys, Seagate and some of the other guys, I think--

anyway, I forgot the name, and said, look, we'll work with you so that you could put a disk drive inside, that was a thicker device, you know, five millimeter, up from 3.3 millimeter, so but that turned out to be too big for cameras and phones. This actually went into an IBM phone. This was an IBM phone, it's called a Simon. It was the brick, IBM shipped that to Bell South. And this was, you know, the funny story on this card is that we didn't sell too many of them because it was really a brick. IBM designed the brick phone for Bell South with the wrong connector on their side. This is a mother connector, so when you would plug it, it would get stuck. So they said-- they came to us and they said, we can't do anything about the phone, can you do anything about the card? So we had an operator spray a lubricant into this thing, so it would slide in and out, a conductive lubricant. Anyway, this (PC card) was too big, so this (CF card) was developed. We developed this (CF card), together with Canon and Kodak, and then this became too big for consumer electronics, so then came the SD card. That (SD card) became too big for cell phones, and that drove us to today, which is the MicroSD which you can barely see, I don't-- but the MicroSD today, which SanDisk makes up to 64 gigabyte, 64 gigabyte storage capacity, with 8 or 16 flash chips sitting inside. The wafer is thinned down, you can see through it, and you stack 8 or 16 flash die, one on top of each other, and a 32 bit microcontroller die, all inside this thing, 64 gigabyte.

Porter: As you're developing all those markets and selling all those increased products at a much reduced price per unit, didn't your production have to go up incredibly high during that period?

Harari: Yes. Our production volume--

Porter: Did you have to keep up with the same amount of money in the till?

Harari: Yes, yes, of course. I mean you're right. You could either sell 100 units or 10,000 units for 1,000 dollars, or a million units for ten dollars. And we prefer the million units for ten dollars, because we always felt that, if we're not going to do it, competition will do it.

Porter: Where were you making all those flash memory products?

Harari: So we actually, we were a fabless company. In a way, we are still. SanDisk still is a fabless company today, except that SanDisk has, so far, invested approximately nine billion dollars, at last count, in manufacturing that captive capacity that SanDisk owns. So we started with AT&T in Allentown, as our manufacturing supplier foundry. They, after a while, we felt that they were not competitive. We moved to Japan, Matsushita. Matsushita was a very good supplier, but after a while, they, too, were unable to meet our cost requirements. From there we went to Korea, LG Semicon. There's another story. They were a good supplier, but up to a point, then we went to Japan, NEC, and then UMC in Taiwan, where we had a good relationship and finally, in the last 11 years, 12 years, with Toshiba, where we have joint venture fabs 50-50, actually, some of the fabs are 40-60 in their favor. So yes, the volume is huge. The volume is huge, the investments are huge, and yeah, you make it the old fashioned way, by selling a lot of volume.

So SanDisk today builds, you know, close to two million cards or units a day, and sells them worldwide, and this would not have been possible if it was not for this cost reduction, you know, the USB flash drive displacing the floppy disk, embedded flash replacing tape in the Sony Walkman, replacing CDs in the MP3 players, now starting to displace DVDs in movies. Your iPhone and iPad, all the tablets that have 128 gigabytes, 64 gigabyte of memory, it's all flash, it's all system flash. There's no way that you can do it without managing the flash.

Porter: While you were evolving all of that growth, were-- and the semiconductor industry and startups, there were scores of other companies starting to produce flash memory products. So you were evolving with growth at the same time, there was a tremendous growth in the number of competitors. Did you set up cooperative arrangements with any of those other companies making flash memories?

Harari: No, the only really important partnership that we had and have, was with Toshiba. Toshiba was very interested in our multilevel cell patents, and knowhow. Multilevel cell was very, very important, and that fact of 30,000 cost reduction, multilevel cell was critical, because it allowed you to store two or three or four bits on every transistor.

Alexy: So you were continuing technology development at a fundamental level, then in the flash market to drive your costs and density.

Harari: Yes, we never stopped and in fact, if anything, I believe that SanDisk is the world's leader in technology development and understanding, together with Toshiba, of fundamental flash scaling, device issues, device scaling and also not just flash, NAND-flash which is the current production workhorse technology, but also 3D memory, vertical NAND, and so on. SanDisk spends a great deal of its operating expenses in R&D, and I was, frankly, personally, very deeply involved on the technology side until the day I retired, because this is my love.

Porter: SanDisk, does it sell flash chips themselves, or is it really the modules with the controller that is the heart of the company?

Harari: From day one, we decided we're not going to sell flash components, stand alone, and we've never changed-- I take that back. We do today now sell wafers, and you can buy a controller from a Taiwanese manufacturer, and buy our wafers and package your own device. But the chips were always intended to be part of a system. There always was an expectation that there was a controller, and the controller knew how the flash worked. So the answer is no, we never sold flash components, and all of our competitors, did. I mean, Intel sold flash components, Toshiba, STmicro and so on. As far as all the companies, the small companies, yes, there were a lot, and in fact, we, our policy, our strategy on licensing, was to license all comers. We did not try to be a Polaroid, you know, you can't step into this thing. On the contrary, we said, you know, for these markets to be very large markets, competition was essential.

Kodak wasn't going to buy flash memory cards if SanDisk was the only supplier. We understood that. Fuji, same thing, so we had to actually bring in competition, but we didn't want to give it away for free, you know, we felt we have developed the market, we have developed the technology, we have the patents, they should pay us because they have the economies of scale that we don't have. And that has worked very well for the company. The smaller companies found that it was just very-- almost impossible for small companies to play on this incredibly rapid technology scaling. Moore's law, think about it, in just-- in 2001, we were at 256 megabit NAND chip, and today we are 64 gigabit NAND chip. That is nine generations of technology, in 11 years. Nine generations in 11 years. For a small startup, it's almost impossible, even for a large company, like Hitachi, that tried to be in this market, AMD that tried to be in this market, it's-- in the space of the last 11 years, we overshot way past the cost of DRAM, cost per bit. Flash, is about a tenth of that of DRAM. And we are, at least for low capacities, we are approaching the cost of our disk drive for low density, you know, hundred gigabytes or so.

Porter: So at the end of the '90s, in 1999, IBM tried to compete with disk drives again, with their one inch microdrive, of which I gave you a sample over there. And it should be noted that there's some-- in the microdrive, as you can notice on the back, you'll see what's inside that drive. There's a cost of all those items involved. There's a couple of-- there's the magnetic disks, there's the heads, there's a motor to make those disks spin, there's another motor to move the actuator, there's probably at least three semiconductors to make those motors under control, and to process the data coming in and out of it, and an incredibly precise group of little metal parts in that mechanism. Now, given all of that, the manufacturer of the drive probably has a cost of about 45 dollars just for the cost of the materials in all of those products, and if he adds a little profit and gross margin on top of that, it's probably a 50 dollar item, at least.

Harari: Yes.

Porter: Now, did you regard that as serious competition? Incidentally, those one inch drives were driven out of the market by flash memory after a couple of years, and I think you had a major part in that, didn't you?

Harari: Yes. So first of all, we worked with IBM to have the ability for that disk drive, microdrive be compatible with our flash. So we did not try to stop IBM coming into that market. We always had confidence that Moore's law would work in our favor. If you look at what you have here, and compare with what we had in our equivalent part, the inside, you can see a controller chip and a memory chip, that's it. No moving parts, none of that. Now, so long as flash chip was 100 dollars or so, very, very expensive, IBM had an advantage, even at 50 dollars cost. But we were able to go way below 50 dollars, of course, so IBM, when they announced their first microdrive, that day our stock dropped by 20 percent, because IBM presented it as the flash killer. They said, we can fit into your camera, we have developed a very low power disk drive, and you get 340 megabytes instead of 8 megabytes for the same price. And initially, they were right. But very, very quickly, we overtook them. They were not unique. We had competition, you

know, the HP Kittyhawk 1.3 inch drive, also tried to displace us. The IOMEGA Clik drive, 40 megabytes on a-- the IBM Millipede that was supposed to be probe technology.

Porter: And they all have that long list of parts, which were made, and therefore had a high cost.

Harari: Yes. And what everybody ignored was the power of Moore's law. And the power of integration and basically, I mean, that controller, which initially was, like 70 dollars, bill of materials, and you know, off the shelf, 70 dollars. Today, the controller that goes in here, is under 20 cents. Under 20 cents, 32 bit micro controller, you know, ARM processor, about 1,000 times faster for, from 70 dollars, to 20 cents. So this is the power of semiconductors, that is that it was like wave after wave of disruption in the marketplace over the last 20 years. The first wave, as I said, is the film. The next wave, the floppy, the next wave, the tape, then the CD, then the DVD, and now coming to the hard disk drive. It's just-- it's just unbelievable, it's a revolutionary force that-- and it's amazing how little respect, frankly, flash gets. The technology inside this, the technology inside, you know, 64 gigabyte, this is a three bit per cell technology. Every transistor will store three bits, 64 gigabyte, it is about 200 billion-- 150 billion transistors, billion transistors. You know, the most advanced Intel microprocessor is, what, 5 billion transistors? So flash now is driving the technology further than anybody else. I mean, flash is at 20 nanometer, 20-- SanDisk has just announced a 19 nanometer NAND flash, to be introduced next year, in production. It will support 128 gigabit on a chip, 16 gigabyte chips, with 16 gigabyte and 16 of these chips, this is 256 gigabyte.

Porter: We have a continuous evolution of mobile applications, with every year, somebody coming up with new mobile applications, something new to fit in your shirt pocket, or something new to put on one of these little book-sized computers, etcetera. Are the applications you have, or your company has still under development, adequate to develop all of these markets?

Harari: The market has always surprised us on the size growing bigger than we thought. I remember around 2005, you know, five, six years ago, when we recognized-- six, seven years ago, we recognized the mobile opportunity and started developing the MicroSD for mobile phones. The-- but I remember sitting in meetings and we were saying, well, you know, you need to have eight gigabyte storage, eight gigabyte storage, in a phone with 100 million phones that would use that. Why would 100 million people need eight gigabyte in a phone? We could not see it coming. But we said, build it, they will come. And today, you know, with the iPhone and the Android phone and so on, and the tablet, of course, it's not just 8 gigabytes, it's more like 32, 64 gigabyte and 100 million is nothing. We didn't see, also, China coming, and India coming, so all the BRIC companies that are now beginning to-- everybody there wants a camera, and everybody there wants a smart phone. So--

Porter: I want to ask a wind up question here. In 2009, Eli, you were awarded the IEEE Noyce Medal. What was your reaction to that kind of award?

Harari: That was the greatest honor that any device guy, any semiconductor guy could hope for. I mean, this was a great honor. Robert Noyce, you know, the inventor of the integrated circuit, really the founder of Fairchild, and founder of Intel, cofounder, the father of semiconductors. I worked at Intel when he was still there. He was already in chairman role, and more the outside guy, so I didn't have a lot of interaction with him. I had a lot more with Gordon Moore, and Andy Grove. But great, great honor, great man. It's just-- it's amazing, it's only-- you know, the integrated circuit was invented in 1959, it's only 52 years ago. Fifty-two years ago, a technology that nothing has ever so radically revolutionized people's lives as the integrated circuit, nothing. I mean, not the combustion engine, I mean, not-- it's just amazing. So-- now interesting, when I got that award, I went to bookstores to learn some more about his life's history, Robert Noyce's, Bob Noyce. I was amazed how few people had even heard his name, which was just shocking. Such a great man who had done so much, so much, and very few people outside-- I mean, it's just amazing. So this is really good that you have a Museum Of Computer History that, at least one place here that--

Porter: Well, I'm a IEEE life member, and I'd like to say congratulations on that award myself.

Harari: Thank you.

Porter: And I'd like to ask you how you feel about your contribution to the industry in establishing that company called, originally for the sun, and later for the sand, your contributions for that company and how you feel about having served all that time until just recently.

Alexy: And your contributions to the semiconductor industry for the technology development.

Harari: Oh, it's a huge privilege, obviously. I mean, I'm the luckiest guy on earth. I was born at the right time, at the right place, and you know, was given opportunity, and I'm an immigrant, to boot, you know, I came to the United States as an immigrant, and I think this is a land of immigrants, and a land of second chances. Certainly I had that. But I didn't do it alone, of course. There were some very, very good people that, together, helped make this happen. It was-- we had some very, very tough times, some very, very hard times. People were always trying to destroy us, to kill us, because we were destroying their business, not that we wanted to. There was a-- last year, 2010, Consumer Electronics Show, I was walking through the booth, and-- with another guy, a SanDisk guy, so when we were going into the Kodak booth, he kind of whispers to me, kind of in jest, he says, "You know, you better not let them know who you are." And I said, I was a little puzzled, I said, "Well, what do you mean?" He said, "Well, you are the man who killed Kodak." And I thought about it, and actually, no, really Kodak killed Kodak. I didn't kill Kodak, nobody killed Kodak, they killed themselves. We worked with Kodak Japan on this compact flash, parts of Kodak understood the value and the power of this technology, and were very strong partners of ours. But the hubris that you get, at some of these companies that have a monopoly, Kodak had 70 percent market share in film. They were 65 percent, or so, gross margin, so they were just milking it, and they just wanted this thing to go away. And they had the technology, Kodak had the CCD technology,

they had the digital imaging, they had everything they needed to displace themselves, but they didn't have the guts to do it, and it was done for them. I mean, we were part of that, we enabled digital photography, but it was not-- we never dreamed that we would-- what had happened. The dramatic-- because silver halide is such a good technology, really, I mean, you have 30 megapixel resolution, you know, perfect images, for $2.99, $3.99 for 24, so-- but you kind of see the extent of the devastation to Kodak. Their market cap today is about one tenth of SanDisk market cap, and I think that we are undervalued. So this is a very unforgiving field. Consumer electronics and mobile computing, you see what's happening with Nokia, going from peak of the hill, you know, Apple comes around, and Android follows, and everything changes so quickly, so quickly. So it's, as I said, incredible privilege to be-- to have had the opportunity as CEO, as founder, as a technologist, to be not just at a front row seat to all of these changes, but actually in the ring. You know, boxing and getting punched from time to time, and having the setbacks from time to time. It was a phenomenal journey, phenomenal journey, and I'm just grateful.

Porter: Well I have had an opportunity to watch the progress of your company and your management over the years, and as you know, you and I have met many times over the years, and all I can say is, personally, congratulations.

Harari: Thank you. Thank you, Jim. Thank you.

Alexy: Same for me.

Harari: Thank you, George. Thank you.

Alexy: Excellent.

Harari: Good, I think we captured it, no? More or less. Good.

Porter: Well, Eli, we had a great discussion here, and thank you very much for coming.

Harari: Thank you. Thank you, Jim, thank you, George.

Alexy: It's been a pleasure.

Harari: Good.

END OF INTERVIEW

